

Creating Micro-Marketing Pricing Strategies Using Supermarket Scanner Data

Alan L. Montgomery

The Wharton School, University of Pennsylvania, Marketing Department, 1400 Steinberg Hall-Dietrich Hall, Philadelphia, Pennsylvania, 19104-6371, alm@wharton.upenn.edu

Abstract

Micro-marketing refers to the customization of marketing mix variables to the store-level. This paper shows how prices can be profitably customized at the store-level, rather than adopting a uniform pricing policy across all stores. Historically, there has been a trend by retailers to consolidate independent stores into large national and regional chains. This move toward consolidation has been driven by the economies of scale associated with these larger operations. However, some of these large chains have lost the adaptability of independent neighborhood stores. Micro-marketing represents an interest on the part of managers to combine the advantages of these large operations with the flexibility of independent neighborhood stores.

A basis for these customized pricing strategies is the result of differences in interbrand competition across stores. These changes in interbrand competition are measured using weekly store-level scanner data at the product level. Obviously, this presents a huge estimation problem, since we wish to measure substitution between each product at a store-level. For a chain with 100 stores and 10 products in a category we would need to estimate over 100,000 parameters. To reliably estimate these individual store differences we phrase our problem in a hierarchical Bayesian framework. Essentially, each store-level parameter can be thought of as a combination of chain-level and random store-specific effects. The improvement in estimating this model comes from exploiting the common chain-level component. In addition, we relate these store-specific changes to demographic and competitive characteristics of the store's trading area, which helps explain why these differences are present.

These estimated differences in price response are in turn used to set store-level prices. To simplify and focus the problem we limit our attention to everyday price changes (i.e., the prices of products that are not advertised). There are many marketing variables that can be adjusted at a store-level (e.g., promotions and product assortments); the reason we concentrate upon everyday pricing is driven by its importance in the marketing mix, that most profits are earned

on products sold at their everyday price, and the amenability of everyday prices to store-level customizations. A limitation of this approach is that it yields only a partial solution to the retailer's global optimization problem.

A challenge for the retailer in implementing micro-marketing pricing strategies is to retain a consistent image while altering prices that adapt to neighborhood differences in demand. Our approach is to search for price changes that leave image unchanged. We argue that a sufficient condition for holding the input to store image constant from everyday pricing is to hold average price and revenues at their current levels. We implement this condition by introducing constraints into the profit maximization problem. Future research into store choice may yield more efficient conditions. A benefit of holding the retailer's image constant is that it does not require costly new information about competitors and promotional activity. Instead, retailers are able to derive these store-level customizations based largely upon scanner data. This is very advantageous since this information is already being collected and is readily available.

Our results indicate that micro-marketing pricing strategies would be profitable and could increase gross profit margins by 4 percent to 10 percent. When these gross profit gains are considered after administrative and operating costs are taken into account, they could increase operating profit margins by 33 percent to 83 percent. These gains come from encouraging consumers through everyday price changes to switch to more profitable bundles of products, and not through overall price changes at the chain-level. These results show that the information contained in the retailer's store-level scanner data is an under-utilized resource. By exploiting this information using newer and more powerful computational techniques managers can better appreciate its value. The implication is that profits could be increased and gains can be made by using this information as the basis for micro-marketing.

(Pricing; Micro-marketing; Segmentation Research; Retailing; Scanner Data; Estimation; Bayesian Hierarchical Models)

1. Introduction

As retailers find their environment becoming increasingly competitive, they have begun to look for new ways to increase their profitability. One method is *micro-marketing*, or the customization of marketing mix elements to the store-level instead of following the same policy for every store in the chain. This move toward micro-marketing reflects a desire by retail managers to retain the cost savings and marketing clout of large national and regional chains while developing more customer-oriented strategies that cater to neighborhood markets, just as small independent retailers have been able to do in the past. This paper shows how profitable micro-marketing strategies can be devised by customizing everyday prices at the store-level using store scanner data.

Creating a micro-marketing policy is a difficult task. Blattberg (1988) writes: "Unless the chain is disciplined, individual stores will vary greatly on pricing and merchandising, destroying overall image and positioning. Therefore, the major cost for a chain is the management effort to reconcile store autonomy with presenting a uniform image to customers."

Our solution to this implementation challenge is to focus the retailer on everyday price changes that will not alter its current image. We concentrate upon everyday pricing both to simplify the problem and focus attention upon the most profitable element of the retailer's pricing strategy. In our dataset, 75 percent of all profits are made on products sold at the everyday price. In addition, everyday pricing is very amenable to micro-marketing since it does not require large advertising outlays, price changes are inexpensive to implement, and, under certain conditions discussed in this paper, everyday price changes will not provoke competitive responses, unlike promotional price changes.

An important facet of our solution is to show how prices can be managed at a store-level while still preserving the retailer's current image. We do this by developing a set of constraints on the pricing problem that are consistent with economic theory, and when they are satisfied will leave image unchanged. An advantage of leaving image unchanged is that it highlights the incremental contribution of micro-marketing and does not confuse these changes with chain-wide

effects. It also counters a criticism by managers to current pricing research (Hoch, Dreze, and Purk 1994), which suggests higher prices and lower sales will yield greater profits. Retailers have been unwilling to adopt the new images implied by these changes since these strategies mean decreased market share, although they have shown a willingness to experiment with store-level strategies.

Currently, many retailers practice a very limited form of micro-marketing pricing, namely zone pricing. Zone pricing is generally implemented by grouping stores into clusters and then proportionately increasing or decreasing all the everyday prices in a cluster. We enrich these simple zone strategies with store-level price changes at a product level. For example, a micro-marketing policy may recommend increasing the price gap between the national brands and the private labels in more affluent neighborhoods to take advantage of low price sensitivity of the more profitable national brands and encourage price-sensitive consumers to buy store brands. This level of pricing management would not be possible with the overall price changes prescribed by a zone strategy. Furthermore, in many instances zone pricing is only used by retailers to respond to competitive conditions, such as the proximity of a warehouse or discount store. However, we account for both competitive and demographic characteristics of the store's neighborhood when estimating price response and deriving new micro-marketing strategies.

The key input to implementing a micro-marketing policy is the measurement of product-level price response for each store. Previous marketing researchers have primarily been interested in market-level or chain-level price response, and not fully considered the problem of measuring demand at a store-level. Gupta, Porter, and Wittink (1993) suggest that data be pooled across stores, while Blattberg and Neslin (1990) and Blattberg and Wisniewski (1989) suggest aggregating store-level data to the zone-pricing level. Although our demand specification is consistent with these previous models, the focus of our problem is on the measurement of demand at a store-level. This presents a formidable estimation problem due to the huge number of parameters. To improve the parameter estimates, Bayesian shrinkage techniques are employed. While

the estimation problem is very interesting by itself (see Montgomery and Rossi 1997), our primary focus is on the pricing implications of these store-level differences in demand.

To gauge the importance of these findings we measure the impact of these store-level differences on expected profits, which has not been previously considered in the marketing literature. This positions our problem in its natural decision-theoretic context, and allows us to assess both the statistical and managerial significance of these results. This contrasts with the usual analysis that stops with the measurement of store-level price elasticities.

The plan of the paper is as follows. In § 2, we discuss the modeling of the retailer's store-level demand function using a hierarchical Bayesian model. Section 3 considers the prediction of profits, and the implementation of micro-marketing pricing strategies. These results are contrasted with the retailer's current zone-pricing strategy in § 4. In § 5, we consider the impact of modeling assumptions on pricing predictions. The implications of these findings are considered in § 6.

2. Modeling Store-Level Demand

To develop micro-marketing pricing strategies we must estimate demand at the store-level. Fortunately,

retailers have a readily available information source that can be used to measure demand: weekly store-level scanner data. While these data are primarily used for inventory and accounting purposes, we show how they can be used to model store-level demand. Our formulation of the micro-marketing problem is driven by this information set.

The dataset used in this paper comes from Dominick's Finer Foods (DFF), a major regional supermarket chain that accounts for 20 percent of supermarket sales in the Chicago area. We model demand for the 11 products ($M = 11$) that make up the refrigerated orange juice (OJ) category. These products are sold in 83 stores ($S = 83$) over a 121-week time span ($T = 121$). Summary statistics computed across the stores for average weekly prices, gross profit margins, and market shares are listed in Table 1.

Each product is modeled as an equation in which its log movement is a function of its own price, the prices of the other products in the category, and its feature and deal status:

$$\ln(q_{its}) = \alpha_{is} + \sum_{j=1}^M \eta_{ijs} p_{jts} + \xi_{is} f_{its} + \psi_{is} d_{its} + e_{its} \quad (2.1)$$

where q_{its} , p_{its} , f_{its} , and d_{its} are movement, price, feature,

Table 1 Descriptive Statistics for Price, Market Share, and Profit Margins Across Stores (Prices are standardized for a 64 oz. unit)

Description	Price for 64 oz.		Market Share				Profit Margin
	Mean	Std Dev	Mean	Std Dev	Minimum	Maximum	
<i>Premium:</i>							
Tropicana Premium 64 oz.	2.87	0.55	16.1%	3.7	9.0%	28.2%	22.6%
Tropicana Premium 96 oz.	3.12	0.39	10.7%	4.1	3.4%	23.1%	25.8%
Floridas Natural 64 oz.	2.86	0.31	4.0%	1.1	2.2%	7.0%	28.9%
<i>National:</i>							
Tropicana 64 oz.	2.27	0.41	15.8%	2.6	10.1%	20.0%	21.5%
Minute Maid 64 oz.	2.24	0.40	16.9%	1.6	13.3%	21.6%	19.4%
Minute Maid 96 oz.	2.68	0.36	5.7%	1.5	2.7%	9.3%	25.3%
Citrus Hill 64 oz.	2.32	0.34	5.1%	0.6	3.7%	6.6%	21.9%
Tree Fresh 64 oz.	2.18	0.29	2.5%	0.7	0.9%	4.0%	27.1%
Florida Gold 64 oz.	2.07	0.41	2.6%	0.6	1.4%	4.2%	27.2%
<i>Store:</i>							
Dominicks 64 oz.	1.74	0.39	13.6%	3.5	5.6%	20.9%	23.8%
Dominicks 128 oz.	1.83	0.32	6.9%	2.5	2.4%	16.4%	27.9%

and deal, respectively. The subscript i represents the product, t denotes week, and s is for store. The feature variable indicates if a product is advertised, and deal indicates any in-store displays like bonus-buy tags or in-store coupons. The vector of error terms, $(e_{1st}, e_{2st}, \dots, e_{Mst})$, follows a multivariate normal distribution with mean $\mathbf{0}$ and covariance matrix Σ_s .

Model (2.1) allows for a rich assortment of inter-brand competition, since it measures the full unrestricted cross-price elasticity matrix. This elasticity matrix is made up of 121 ($= 11 \times 11$) own- and cross-price terms for each store, or a total of 10,043 ($= 121 \times 83$) terms at the chain-level. Alternative specifications with cross-promotional effects and regular versus promotional prices were tried and could be incorporated in our method, but the specification of (2.1) was favored by the Schwarz Information Criterion (SIC) and out-of-sample predictive validation tests at the store-level. These specification tests indicate the primary empirical problem is whether to pool the data across the stores or use individual store models. Residual diagnostic tests show little serial correlation, which supports the adequacy of this model.

We do not include the prices of other similar DFF categories like frozen juices, since pricing experiments show no significant substitution between the frozen and refrigerated juice categories (Hoch and Purk 1993). In addition, Model (2.1) does not include prices of competing retailers, since we were unable to obtain these data. Previous research suggests that their omission will be a minor issue since within-store substitution is much higher than cross-store substitution. Store-level research by Kumar and Leone (1988) showed that within-store substitution rates were two to three times greater than substitution across stores induced by promotions. Their results should be stronger than those for our category, since their category—disposable diapers—is generally used as a loss-leader and has a high per unit item price, while ours is neither. However, since prices of competing retailers are not included we are careful to only consider price changes that will not induce competitive reactions (see § 3).

2.1. Improving Individual Store Model Estimates Using Hierarchical Bayesian Models

A natural approach to estimating (2.1) is to let each store have its own model. However, these individually

fit store estimates have large standard errors that suggest some of the observed store to store differences may be the result of estimation error. One approach to improving these estimates is to pool across the stores (see Wittink et al. 1988). Pooling results in very stable parameter estimates, but these gains are made by ignoring store-level differences, which impedes the development of a micro-marketing strategy. In addition, these pooled models are strongly rejected in favor of store-level models by likelihood ratio tests ($F = 14.2$, the 99th percentile for this F test is 1.03) and increase the out-of-sample mean-squared error (MSE) over individual store models by 14 percent (see § 5 and Table 7).

To solve these empirical problems we frame our model in a hierarchical Bayesian context (Lindley and Smith 1972, Smith 1973, Gelfand and Smith 1990, Gelfand et al. 1990). This framework improves the individual LS estimates by borrowing information across the stores and shrinking the estimates toward one another. Under certain prior assumptions, our model will produce estimates identical to either pooled or individual store-level models. However, we choose a more moderate prior that will result in estimates that fall between those of pooled and individual store models. To support our specification we find that out-of-sample MSE is reduced by 20 percent over individual store models and a 30 percent reduction compared with a pooled model (see § 5).

Certainly there are other methods of estimating store-level models that can help stabilize store-level parameter estimates. For example, a specific market structure could be assumed, such as a macro-logit model, constraints on the cross-elasticity matrix could be imposed (Allenby 1989), or only selected parameters could be allowed to vary across stores (Hoch et al. 1995). These methods are not used, since we are especially interested in allowing for market structure changes across stores and do not wish to impose fixed structures. Our goal is to let the data drive the results as much as possible. To the extent that demand is difficult to estimate (e.g., high standard errors), the uncertainty in the parameter estimates is incorporated into the posterior of the profit function (see appendix).

2.2. The Hierarchical Bayes Model

Model (2.1) can be rewritten in matrix form:

$$\begin{bmatrix} \ln(q_{11s}) \\ \ln(q_{21s}) \\ \vdots \\ \ln(q_{M1s}) \end{bmatrix} = \begin{bmatrix} \alpha_{1s} \\ \alpha_{2s} \\ \vdots \\ \alpha_{Ms} \end{bmatrix} + \begin{bmatrix} \eta_{11,s} & \eta_{12,s} & \cdots & \eta_{1M,s} \\ \eta_{21,s} & \eta_{22,s} & \cdots & \eta_{2M,s} \\ \vdots & \vdots & \ddots & \vdots \\ \eta_{M1,s} & \eta_{M2,s} & \cdots & \eta_{MM,s} \end{bmatrix} \begin{bmatrix} p_{11s} \\ p_{21s} \\ \vdots \\ p_{M1s} \end{bmatrix} \\
+ \begin{bmatrix} \xi_{1s} \\ \xi_{2s} \\ \vdots \\ \xi_{Ms} \end{bmatrix} \begin{bmatrix} f_{11s} \\ f_{21s} \\ \vdots \\ f_{M1s} \end{bmatrix} \\
+ \begin{bmatrix} \psi_{1s} \\ \psi_{2s} \\ \vdots \\ \psi_{Ms} \end{bmatrix} \begin{bmatrix} d_{11s} \\ d_{21s} \\ \vdots \\ d_{M1s} \end{bmatrix} \\
+ \begin{bmatrix} e_{11s} \\ e_{21s} \\ \vdots \\ e_{M1s} \end{bmatrix} \quad (2.2)$$

or equivalently in vector notation as:

$$\mathbf{q}_{1s} = \boldsymbol{\alpha}_s + H_s \mathbf{p}_{1s} + \text{diag}(\boldsymbol{\xi}_s) \mathbf{f}_{1s} + \text{diag}(\boldsymbol{\psi}_s) \mathbf{d}_{1s} + \mathbf{e}_{1s}, \quad \mathbf{e}_{1s} \sim N(\mathbf{0}, \boldsymbol{\Sigma}_s). \quad (2.3)$$

The parameters from a store's demand system (β_s) can be stacked into a single vector:

$$\boldsymbol{\beta}'_s = [\boldsymbol{\alpha}'_s \text{vec}(H_s)' \boldsymbol{\xi}'_s \boldsymbol{\psi}'_s]. \quad (2.4)$$

Our hierarchical model treats β_s as a draw from a latent distribution, known as the hyper-distribution:

$$\boldsymbol{\beta}_s = \mathbf{Z}_s \boldsymbol{\theta} + \mathbf{u}_s, \quad \mathbf{u}_s \sim N(\mathbf{0}, V_{\boldsymbol{\beta}}). \quad (2.5)$$

The model for the j th element of the β_s vector can be rewritten as:

$$\beta_{js} = \bar{\beta}_j + \gamma_j' \mathbf{d}_s + u_{js}, \quad (2.6)$$

where \mathbf{d}_s is the vector of demographic and competitive variables for the trading area of store s . Note that while the prior depends only upon the demographic information, the posterior distribution of β_{js} will depend upon both the demographics and store sales data. Intuitively, our posterior estimates can be thought of as a weighted average of the usual LS estimates and the pooled estimates with demographic interactions, at least approximately.

Table 2 defines and describes the demographic and competitive variables (\mathbf{d}_s) from (2.6). To be able to interpret the barred constants as chain-wide averages, the \mathbf{d}_s vectors are standardized with zero means ($\boldsymbol{\Sigma}_d = \mathbf{0}$). Also instead of allowing each of the 154 parameters a separate set of demographic parameters (γ_j), we use the same demographic effects for five types of variables: constants (γ_a), own-price (γ_o), cross-price (γ_c), feature (γ_f), and display (γ_d) terms. This introduction of demographic effects to estimate store-level price response has not been previously considered.¹

To complete the Bayesian specification of our hierarchical model we place natural conjugate priors on the error covariance matrix $\boldsymbol{\Sigma}_s$ given in (2.3):

$$\boldsymbol{\Sigma}_s^{-1} \sim \text{Wishart}(\nu_s, \bar{V}_s^{-1}) \quad (2.7)$$

¹We refer the interested reader to Rossi, McCulloch, and Allenby (1996) for an additional study that employs demographic variables in a hierarchical Bayesian framework to explain consumer heterogeneity using a household-level probit model.

Table 2 Descriptive Statistics for Demographic/Competitive Variables of the Store's Trading Area Across the Chain's 83 Stores

Variable	Description	Average	Std. Dev.	Minimum	Maximum
Elderly	% of population over age 60	0.173	0.062	0.058	0.307
Educ	% of population with a college degree	0.226	0.111	0.050	0.528
Ethnic	% of population that is black or Hispanic	0.155	0.188	0.024	0.996
Income	Log of median income	10.618	0.283	9.867	11.236
Fam_Size	% of households with five or more members	0.116	0.030	0.014	0.216
Work_Wom	% of women who work	0.359	0.053	0.244	0.472
House_Val	% of homes with a value greater than \$150,000	0.345	0.241	0.003	0.917
Ware_Dis	Distance (miles) to nearest warehouse	6.150	3.790	0.132	17.856
Ware_Val	Ratio of DFF store sales to nearest warehouse	1.321	0.493	0.500	3.273
Super_Dis	Average distance (miles) to nearest five supermarkets	2.118	0.738	0.773	4.108
Super_Val	Ratio of DFF store sales to average of nearest five supermarkets	0.452	0.206	0.096	1.114

and prior distributions on the mean and covariance matrix of the second stage given in (2.5):

$$\theta = [\beta' \gamma_a' \gamma_o' \gamma_c' \gamma_f' \gamma_d']' \sim N(\theta, V_\theta), \quad (2.8)$$

$$V_\beta^{-1} \sim \text{Wishart}(\nu_\beta, \bar{V}_\beta^{-1}).$$

The motivation for representing V_β with a-prior instead of specifying it directly is to allow for some uncertainty in the amount of commonalities across stores. For a discussion of the remaining priors on the parameters of the hyper-distribution and the estimation of this model, see the appendix.

The general framework we have proposed here is similar to Blattberg and George (1991) but differs in several key ways. These differences largely reflect the divergence of interest between Blattberg and George (1991) and this paper. Their primary interest was to shrink the brand estimates toward a central tendency to avoid any nonsensical estimates, i.e., positive own-price elasticities and negative promotional estimates.

However, our interest in the estimation of store-level demand estimates is to use these estimates as an input to the profit function (which we will discuss in the following section). Our dataset is much richer, so we do not concern ourselves with shrinkage across brands but instead focus on pulling the store-level estimates toward one another. To help explain cross-store variation we introduce a demographic predictor, which results in our cross-store estimates being shrunk toward a regression line and not a single point. Our model also differs in its treatment of demand as a system of equations with correlated errors and the parameterization of the prior. Finally, our estimates show strong improvements in out-of-sample prediction, while theirs did not (see § 5).

Empirical Results. The estimates of the posterior mean (θ) and standard deviation of the hyper-distribution (V_β) are listed in Table 3. The price sensitivity parameters are multiplied by the average price

Table 3 Posterior Mean (θ) and Standard Deviation (V_β) of the Hyper-Distribution

	Cross Price Elasticity Matrix evaluated at average prices ($H \text{ diag}(\bar{p})$)										Other Parameters			
	TropP 64	TropP 96	FINat 64	TropR 64	MM 64	MM 96	CHill 64	TFrsh 64	FGold 64	Dom 64	Dom 128	Const α	Deal ψ	Feat ξ
TropP64	-2.93	0.15	0.18	0.17	0.02	-0.03	0.11	0.29	0.14	0.05	0.43	10.65	-0.06	0.57
	1.00	0.75	0.65	0.39	0.25	0.75	0.43	0.46	0.31	0.25	0.79	1.18	0.12	0.18
TropP96	0.38	-1.84	0.03	-0.01	0.02	0.15	0.03	0.19	-0.04	0.08	0.00	9.64	0.10	0.29
	0.22	0.96	0.30	0.22	0.17	0.52	0.28	0.34	0.21	0.16	0.65	1.07	0.11	0.16
FINat64	0.89	0.32	-3.72	0.36	0.14	-0.19	-0.04	0.02	-0.24	0.10	0.72	9.29	0.02	0.57
	0.43	0.98	1.22	0.42	0.35	0.94	0.65	1.21	0.45	0.39	1.70	1.84	0.17	0.23
TropR64	1.24	0.92	-0.13	-3.82	0.96	1.60	0.63	0.04	-0.06	0.08	-1.41	8.69	0.09	0.71
	0.76	0.90	0.77	0.85	0.67	1.08	0.88	0.65	0.58	0.53	1.43	1.90	0.20	0.26
MM64	0.53	0.33	0.84	0.65	-2.88	-0.54	0.38	-0.03	0.12	0.55	0.05	8.94	0.06	0.80
	0.36	0.66	0.47	0.34	0.83	0.75	0.63	0.56	0.41	0.30	0.70	1.20	0.16	0.28
MM96	0.04	0.45	0.08	0.01	0.18	-2.41	0.12	0.14	0.00	0.02	0.15	9.43	0.02	0.22
	0.25	0.51	0.39	0.29	0.29	0.98	0.40	0.40	0.27	0.17	0.71	1.07	0.09	0.15
CHill64	0.52	0.25	0.09	0.17	0.74	0.28	-3.52	0.05	0.31	-0.14	0.55	8.80	-0.04	0.96
	0.35	1.13	0.64	0.41	0.50	1.03	1.29	0.52	0.38	0.30	1.19	1.54	0.19	0.36
TFrsh64	0.05	0.26	0.22	0.20	0.24	-0.02	0.34	-3.04	0.15	0.30	0.87	7.91	-0.06	0.47
	0.32	0.69	0.64	0.34	0.28	0.76	0.58	1.82	0.52	0.27	0.92	1.67	0.15	0.22
FGold64	0.21	0.25	-0.10	0.73	0.86	0.14	0.50	0.35	-3.90	0.35	0.74	7.28	0.00	1.09
	0.75	1.45	1.26	0.70	0.78	1.43	1.19	1.12	1.35	0.47	1.43	2.69	0.30	0.38
Dom64	0.22	-1.58	0.30	1.15	1.13	0.02	-0.06	-0.15	0.20	-2.68	-0.26	10.69	0.03	0.86
	0.37	1.36	0.78	0.82	0.75	1.10	1.01	0.92	0.48	0.82	1.29	2.07	0.43	0.38
Dom128	-0.13	-0.07	-0.01	0.06	-0.04	0.39	0.09	-0.04	-0.06	0.14	-1.36	9.78	0.10	0.39
	0.20	0.57	0.33	0.27	0.32	0.60	0.39	0.45	0.29	0.24	1.26	1.23	0.18	0.22

to yield the average price elasticity matrix.² The standard deviation of the random store fluctuations are the square root of the diagonal elements of the posterior mean of V_{β} . (Note V_{β} gives the standard deviation of the draws from the hyper-distribution and not the standard error of the estimates.) We can observe quite large values of V_{β} , which demonstrates that store to store variability in price response exists.

To illustrate this table consider Minute Maid 64 oz. If we were to make an inference about a new store without any demographic information, we would expect its average own-price elasticity to be -2.88 . The standard deviation of .83 indicates there is quite a bit of variability across stores. (To evaluate the total store to store variation in price response we need to consider both the random variation induced by V_{β} and the variation induced by the demographics, which is discussed next.) When this brand is featured we would expect an 80 percent increase in movement, while an in-store display would only increase movement by 6 percent. Notice that most brands are quite sensitive to feature promotions, but show little effect of in-store deal promotions due to the difficulty of creating displays in a refrigerator case.

The relationship between store to store changes in demand and the demographic and competitive variables of a store's trading area are provided in Table 4. To illustrate these relationships consider the effects of *house_val* on own-price sensitivities. A store with 75 percent of the houses in its trading area with values of \$150,000 or more (the average is 35 percent from Table 2) would result in all own-price sensitivities increasing by 8.5 $(=.75 - .35) \times 21.25$ over the average, or an increase in own-price elasticities of .30. This means more affluent areas are less

²For the most part the cross-price terms are positive, which indicates substitution. However, the negative cross terms, which imply complementary behavior, may be explained by category expansion—price reductions of certain brands may attract more consumers into the category. This may induce higher sales of products that would otherwise appear to be independent or weak substitutes. Alternately, there may be a spurious correlation in the pricing relationships due to the infrequent promotions of the larger sizes. We have not attempted to introduce additional prior information, i.e., sign constraints or cross-brand shrinkage, to change the sign or magnitude of these coefficients, instead we incorporate any parameter estimation error into the posterior of profits.

Table 4 Posterior Mean (0) and Standard Deviation of Demographic and Competitive Effects upon Price and Feature Sensitivity

	Constant	Own-Price	Cross-Price	Deal	Feature
Elderly	1.13	-9.77	1.88	-0.06	0.43
	1.46	30.81	3.85	0.17	0.24
Educ	1.17	-15.97	-0.15	0.05	-0.07
	1.17	24.24	2.92	0.13	0.18
Ethnic	0.89	-11.60	0.50	-0.02	0.12
	0.43	8.69	1.14	0.05	0.07
Income (log)	-0.32	-0.55	0.46	-0.02	0.02
	0.37	7.72	0.95	0.04	0.06
Fam_Size	0.59	-47.25	-2.00	0.12	0.78
	2.57	53.16	6.73	0.30	0.41
Work_wom	-1.01	13.93	0.06	0.15	-0.19
	1.64	35.47	4.38	0.20	0.27
House_Val	0.43	21.25	-1.81	0.00	0.02
	0.47	9.97	1.24	0.06	0.08
Ware_Dis (miles)	-0.01	0.21	-0.05	0.00	0.00
	0.02	0.34	0.04	0.00	0.00
Ware_Vol	-0.01	-0.77	-0.28	0.01	0.01
	0.11	2.45	0.30	0.01	0.02
Super_Dis (miles)	0.15	-0.08	-0.19	0.00	-0.03
	0.07	1.57	0.20	0.01	0.01
Super_Vol	-0.25	-1.69	-0.71	-0.02	0.02
	0.31	6.48	0.82	0.04	0.05
R^2	.14	.22	.02	.05	.25

price sensitive (i.e., more positive). The R^2 measure shows that these demographic effects are moderately predictive for the feature, own-price, and constant terms.

It is not necessary to understand why price variation occurs in order to develop a pricing strategy that adapts to these differences. Nevertheless, these results show that there are some interesting effects that explain why price sensitivity varies across stores (i.e., more price sensitive stores are located in areas with higher own-price sensitivities). The fact that these demographic variables are only mildly predictive illustrates the importance of using historical data to estimate price response and not simply using demographic data alone. However, it should be noted that stronger priors, i.e., more shrinkage, results in stronger demographic relationships (Montgomery 1994).

To depict the shrinkage effects of our hierarchical Bayesian model, we overlay the estimates of the

own-price sensitivity for Minute Maid 64 oz. using individual LS store models, a pooled model, and the Bayes estimates for our hierarchical model in Figure 1. The vertical axis is own-price sensitivity, and the horizontal axis is the prediction of the parameter estimate using only demographic information. (See Table 4 for the dimensions of this demographic predictor.) The open diamonds represent the LS estimates and the solid diamonds represent the estimates from the Bayesian model. To picture the shrinkage of the LS estimates toward a central tendency, which in this case is a demographic predictor denoted by a solid line, we draw a dotted line with an arrow pointing from the LS estimate toward the Bayesian estimate. Notice that most Bayes estimates fall between the pooled estimates denoted by a dashed line and individual LS estimates, with the amount of shrinkage depending upon the reliability of the estimate.



3. Measuring Profits from Everyday Micro-Marketing Pricing Strategies

The previous section showed that there are store-level differences in price response; in this section we translate these differences into their impact on profits. The retailer's objective is the maximization of the present value of profits across all stores. If we ignore the discount rate and assume that the retailer has identical profit expectations in future years, we can consider expected chain profits for a given year:

$$E[\Pi] = \sum_s E[\Pi_s] = \sum_s \sum_{t=1}^{52} E[\pi_{t,s}], \quad (3.1)$$

where Π denotes the chain profits, and Π_s are the profits for store s . The retailer's weekly expected gross

Figure 1 Bayes Shrinkage Effects on Individual LS Model Estimates for Minute Maid 64 oz.



profit for store s during week t , denoted as $E[\pi_{its}]$, is implied by Model (2.1):

$$E[\pi_{its}] = \sum_i (p_{its} - c_{it}) E[q_{its}] \text{ and}$$

$$E[q_{its}] = \exp \left\{ \alpha_{is} + \sum_j \eta_{ijs} p_{jts} + \zeta_{is} f_{its} + \psi_{is} d_{its} + \frac{1}{2} \sigma_{is}^2 \right\}, \quad (3.2)$$

where c_{it} denotes the retailer's cost of product i in week t , which we assume is the same for all stores, and σ_{is}^2 is the i th diagonal element of Σ_s .

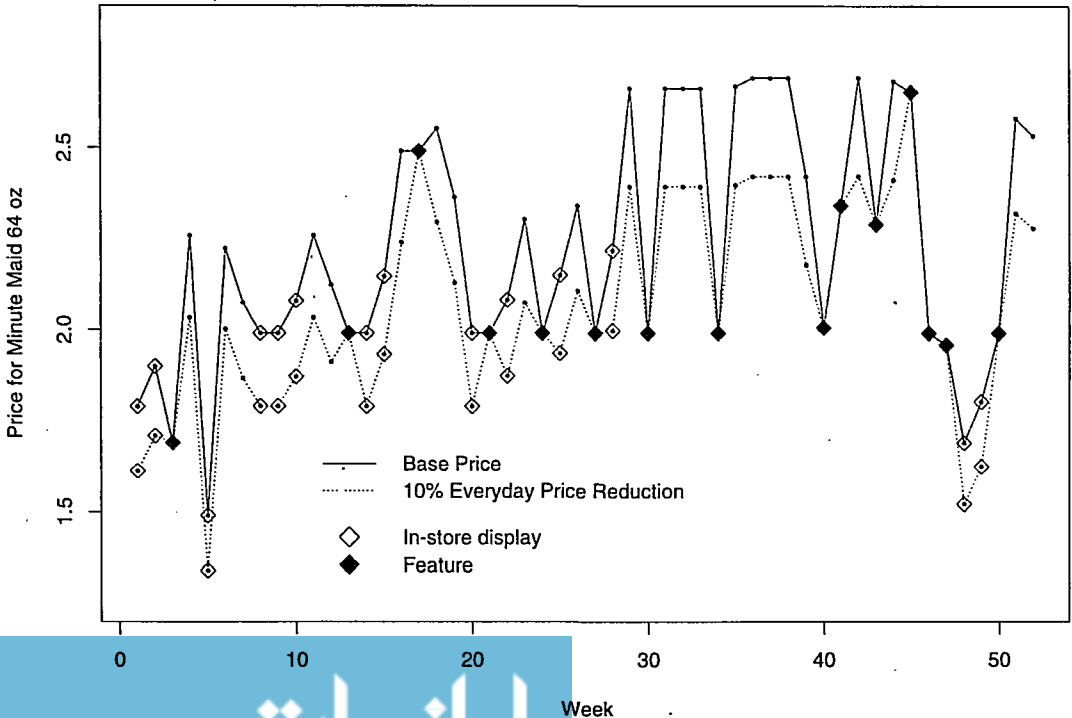
As discussed in the introduction we only consider changes to everyday prices. This will yield a partial solution to the retailer's global profit optimization problem. The everyday price for any given store and week can be thought of as a change from a base pricing policy:

$$p_{its} = \begin{cases} p_{it}^* & \text{if } f_{its} = 1 \\ p_{it}^* \rho_{is} & \text{otherwise,} \end{cases} \quad (3.3)$$

where ρ_{is} refers to a price multiplier for product i in store s , and p_{it}^* is DFF's base pricing strategy. Notice that prices of featured promotions will be the same throughout the chain. The base prices, features, and costs are the actual ones used by the retailer during a one-year period.

To illustrate how our proposed changes will affect price, we plot the base prices of our 52-week period for Minute Maid 64 oz. OJ ($i = 5$) in Figure 2. The base pricing strategy is depicted by a solid line with featured prices denoted with a solid diamond. The variation through time of this base pricing strategy is induced by changes in wholesale costs and promotional offers by the manufacturer. Observe that most large

Figure 2 The Effects of a 10 Percent Reduction in Everyday Prices ($\rho_{s,t} = .9$).



price reductions are accompanied by features, and the minor ones by in-store displays. The retailer's micro-marketing decision is how to customize this brand's everyday price for each store, which is denoted by p_{5s} in (3.3). If everyday prices for this product, in store 1, are decreased by 10 percent, we set $p_{5,1} = .9$. The resulting price series is denoted by a dashed line in Figure 2. Notice that while everyday prices and prices during weeks with in-store displays are reduced by 10 percent, prices during featured weeks are left unchanged.³ We remind the reader that each product has a different base pricing strategy.

We assume that the retailer has decided upon the base strategy (p_i^*) and promotional policy (f_{its}), and must only decide upon the store-level deviations (ρ_{is}) away from these base levels. In the full pricing problem in (3.1) the retailer must make pricing decisions for each product, every week, and in all stores ($11 \text{ products} \times 52 \text{ weeks} \times 83 \text{ stores} = 47,476$ pricing decisions). By separating the pricing decisions as in (3.3), the retailer must now decide upon a base pricing strategy ($11 \text{ products} \times 52 \text{ weeks} = 572$ pricing decisions) and how to modify this strategy for each product in every store ($11 \text{ products} \times 83 \text{ stores} = 913$ pricing decisions).

The decomposition of the pricing problem in this way reflects current management practices (Hoch and Purk 1993) in which zone and everyday pricing decisions are considered separately from base pricing decisions. Given the huge number of pricing decisions that must be made, supermarkets generally choose target or base price levels based upon the current costs and competitive actions of other supermarket retailers and then assign stores to specific zones

³One limitation of our proposed micro-marketing pricing is that "%-off" and absolute dollar amounts for a given product cannot be advertised together. For example, suppose the regular price in one store is \$2.20 and in a second store it is \$2.30. Certainly, the chain's feature advertisements could read "10%-off regular price" or "sale price of \$1.99." However, the chain cannot advertise "10%-off regular price at \$1.99," since this would not be true for all stores. We have operationalized featured prices as absolute dollar amounts (e.g., "sale price of \$1.99"), since we do not have information about the advertising copy. However, our pricing framework can easily be extended to deal with both advertisement methods (i.e., some products can be advertised at a "%-off" or absolute dollar amount) if this information is available.

based upon location. Weekly deviations from this everyday strategy are made because of promotional offers from the manufacturer to the retailer (Armstrong 1991). The development of promotional schedules tends to be a highly competitive activity and is strongly influenced by manufacturer promotions and forward-buying decisions.

The benefit of our approach to setting micro-marketing pricing policies using everyday prices is that it focuses attention on the micro-marketing aspect of pricing, i.e., ρ_{is} and not p_i^* . In addition, it does not require the additional competitive and cost information that is needed to set the base pricing strategy (p_i^*). Instead, this information is indirectly inherited through the base strategy. It is important to recognize that decisions about an appropriate base pricing strategy and promotions are very difficult ones all by themselves, since they include timing decisions (promotional calendars, expected competitive responses, forward buying, etc.) and setting relative prices between products (e.g., the price gap between premium and store brands, subsequent effects on manufacturer relationships, etc.). The disadvantage of this decomposition is the loss of potential profits due to interactions between base pricing policies and micro-marketing, which means the potential profit gains from micro-marketing will be understated.

3.1. The Costs of Everyday Micro-Marketing Pricing Strategies

The physical costs associated with everyday micro-marketing price changes are small. Advertising promotional offers (i.e., featured price discounts) is a very expensive proposition since it requires large outlays for newspaper, direct mail, and television advertising. However, everyday prices are observed in-store and are not advertised. Therefore, changes in everyday prices would not incur additional advertising outlays. In addition, the physical costs of actually changing a price are very small (Goodstein 1994, Garland 1992, Cutler and Rowe 1990). Prices are input at a central location using a computerized inventory and pricing control system. These prices are then electronically distributed to each store, where shelf tags are printed and attached for those items whose prices have changed. (DFF does not tag individual items, so only

shelf tags need to be changed.) There would be no additional cost from printing the shelf price tags due to our micro-marketing price changes, but beyond this start-up cost there would be no additional physical costs incurred.

Besides the physical costs there are also managerial and analytical costs to the retailer. However, if scanner and decision support systems are already in-place—as is the case for DFF—the incremental costs to micro-marketing will be fairly low (Blattberg 1988). An upper bound on these managerial costs can be inferred from the current returns to zone-pricing strategies, which also require store-level decisions. (In § 4, we estimate these returns to be around \$13,000 for the OJ category.) Since the techniques used in this paper are readily available and can be implemented on a micro-computer, the analytical costs are small. The incremental cost of maintaining a micro-marketing pricing strategy over a zone-pricing strategy is negligible, since our pricing scheme will not increase the frequency of week to week price changes. A method for avoiding a final and potentially major cost of micro-marketing due to responses of consumers and competing supermarkets is considered next.

3.2. Micro-Marketing Pricing Strategies

Our basic presumption is that the retailer is currently at a profitable and viable store image, and what is desired from a micro-marketing policy is the preservation of this image. It is an unsettled matter in the marketing literature what precisely determines store image, or store traffic and choice (Bell 1995, Bodapti 1996, Dreze 1995, Walters 1991, Walters and MacKenzie 1988). However, we do know everyday prices, promotions, price format, advertising, loss-leaders, couponing, etc. are inputs to image. One way to leave store image unchanged is to hold these inputs constant. Since our proposed micro-marketing policies will only affect everyday prices, we concentrate upon holding the everyday price input to store image constant. This does not mean retailers must naively match every price of a competing store. Consumers have very poor recall of specific prices (Dickson and Sawyer 1990); therefore, it would appear that interstore competition must occur through general price levels.

We propose two conditions that we argue are

sufficient to preserve a retailer's current image. The first is that average prices should be held at their current levels:

$$\sum_i \sum_j w_{ijs} p_{ijs} = \bar{p}_s, \quad (3.4)$$

where \bar{p}_s is the current average price from store s and w_{ijs} is the market share and is defined as $w_{ijs} = p_{ijs} E[q_{ijs}] / x_{is}$ and $x_{is} = \sum_j p_{ijs} E[q_{ijs}]$. The second constraint is to hold sales revenue constant:

$$\sum_i \sum_j p_{ijs} E[q_{ijs}] = x_s, \quad (3.5)$$

where x_s are the current sales in store s . Notice that these constraints are proposed at the store-level and not the chain-level. Essentially our micro-marketing strategies allow the retailer to increase price variation within stores, but average price levels for each store are held constant. These restrictions are quite severe since they mean that any everyday price increases must be offset by price decreases. Therefore, profit increases result from encouraging consumers through everyday price changes to switch to bundles of products with the same dollar value but higher total profits.

These constraints have a great deal of intuitive meaning to retailers. The average price constraint can be thought of as trying to hold the price image of the store at their current level. The revenue constraint is essentially an attempt to hold market share constant. Assuming that total Chicagoland grocery sales are approximately fixed, if DFF were to increase its revenue, the revenues of competing retailers would necessarily decrease—provoking defensive competitive responses. By holding revenue constant, this implies DFF's market share also remains constant—minimizing the likelihood that competitors would respond. If DFF's current market share is 20 percent before a micro-marketing strategy, its market share after implementing a micro-marketing strategy will remain at 20 percent.

Economic theory provides a rationale for how these constraints hold price image constant. An implicit assumption in our demand Model (2.1) is that the utility for this category in a given store is weakly separable from all other products:

$$utility = f(v(p), v(p_0), v(p_c)), \quad (3.6)$$

where p is the vector of prices for refrigerated orange juices, p_0 is the vector of other prices in the store, and p_c is the vector of prices in competing supermarkets. Weak separability implies that consumer shopping decisions can be described in a hierarchical budgeting process: First, consumers decide upon category expenditures by allocating their budget based upon category price indices and then make individual product decisions within a category subject to these initial category expenditure allocations. (See Chapter 5 in Deaton and Muellbauer 1980 for an introduction and further references.)

Our constraints guarantee that this initial category expenditure allocation will not be affected since average prices and revenues are held constant. Upon evaluation of the indirect utility function we find:

$$f(v(p), v(p_0), v(p_c)) \approx utility^* \quad \forall p \quad (3.7)$$

that satisfy (3.4) and (3.5), where $utility^*$ is constant and denotes the initial utility before any micro-marketing changes. Equation (3.7) holds only approximately since it uses an average price index. To achieve equality, the price index must be a function of prices and current utility (see Chapter 7 in Deaton and Muellbauer 1980).

Equation (3.7) means that any price changes that follow these constraints will not change aggregate consumer utility. Subsequently, competing retailers will not have an incentive to react since the initial expenditure allocations made by consumers to their stores, as prescribed in our hierarchical budgeting process, will remain unchanged. Consequently, all pricing strategies that satisfy these constraints will result in similar competitive environments. Presumably, demand and supply are currently at a stable equilibrium, which implies that our proposed micro-marketing pricing strategy will inherit this current stable equilibrium and preserve the competitive environment. In practice, we believe that (3.7) will hold only approximately due to the complexities of aggregating across consumers, hence this approximation may be inaccurate under some configurations of heterogeneity in

compensatory decision rules where a consumer considers only the minimum price).

3.3. Optimal Product-Level Micro-Marketing Strategies

Following the discussion in the previous subsection, we can now estimate optimal expected profits given in (3.1) by choosing ρ_{is} subject to the constraints given in (3.4) and (3.5) (see the appendix for a discussion of estimating profits). These optimal values are computed numerically using the E04VCF subroutine in the NAG statistical software library (NAG 1991) since an analytical solution for this model is not known. As a basis for comparison, we use a uniform chain pricing strategy, i.e., one that keeps the prices of each product the same in all stores. This can be represented in (3.3) by letting $\rho_{is} = 1$ for all i and s . A uniform chain strategy results in expected profits of about \$3.3mm per year, an average chain price of \$2.23 per 64 oz. carton, and total chain revenues of \$13.9mm. Table 5 lists the posterior mean and standard deviation of expected profits.

Before moving to a store-level strategy, it is worthwhile to ask the question of whether the retailer can improve its current uniform chain pricing policy and yet still satisfy the store constraints on average price and revenue given in (3.4) and (3.5). To do this, we find the price multipliers ($\bar{\rho}_i$) for each product i , where $\bar{\rho}_i = \rho_{is}$ for $s = 1, \dots, S$, that maximize chain profits. This means prices for a given product are the same across the stores. The solution to this problem yields a .4 percent increase in expected profits (line 1 in Table 5).⁴ To illustrate this solution, we plot the price change ($\bar{\rho}_i$) in Figure 3 and denote them with a diamond. For example, under this optimal uniform chain strategy the everyday price of Minute Maid 64 oz. would increase by 1.4 percent over current levels. Notice that all suggested price changes are small (less than 5 ¢ per carton). While this new chain-wide pricing strategy is better than the current one, the fact that it is close to the current one of the base price zone tends to support our argument that these constraints approximate the

⁴To avoid numerical difficulties we relax these revenue constraints so that only chain revenue is unchanged and not the entire vector of store revenues. Therefore, this .4 percent profit increase is overstated.

Table 5 Expected Gross Profit Changes Under Various Constrained Pricing Strategies

Description of Pricing Strategy	Relative to a Uniform Chain Pricing Strategy			
	Expected Profits	Expected Increase	% Change in Expected Profits	Prob[Expected Increase >0]
Uniform Prices across all Stores	\$3,330,900 (11,900)			
1 Optimal Uniform Strategy	\$3,344,100 (11,600)	+ \$13,200 (1,700)	+ .4% (.1)	1.00
2 Optimal Micro-Marketing Strategy	\$3,459,000 (19,400)	+ \$128,100 (18,400)	+ 3.9% (.6)	1.00
3 Optimal Micro-Marketing Strategy with constraints at the Chain-level	\$3,481,600 (20,900)	+ \$150,700 (20,200)	+ 4.5% (.6)	1.00

Note. The standard deviations of the posterior are given in parentheses below the posterior means.

current solution found with the retailer's unobserved full-information set.

We now turn to the micro-marketing optimization problem. Here we must choose a total of 913 different price multipliers for each store and brand (ρ_{isr} for $i = 1, \dots, 11$ and $s = 1, \dots, 83$) that maximize chain profits subject to the constraints in (3.4) and (3.5) for each store.⁵ Since there are no interactions between stores, this maximization problem reduces to optimizing the expected profits in each store, $E[\Pi_s]$, by choosing $\rho_{1sr}, \rho_{2sr}, \dots, \rho_{11,s}$ subject to the constraints that average price (\bar{p}_s) and revenue (x_s) are held at their current values. (Alternately, we could view this optimization problem as relaxing the uniform chain pricing constraint, $\bar{p}_i = \rho_{isr}$, which is enforced in the previous problem, line 1 of Table 5.) The solution to this problem yields a 3.9 percent increase in expected profits, as reported in line 2 of Table 5. This 3.5 percent profit increase over the optimal uniform chain policy represents the incremental contribution of micro-marketing, since there are no chain-wide differences in this strategy.

⁵We also restrict any price changes to within ± 30 percent to guarantee that all profit solutions are well behaved. These boundary conditions are not binding in 96 percent of the pricing changes, and for the most part can be ignored. The conditions that are important in finding these optimal profit solutions are the revenue and average price constraints.

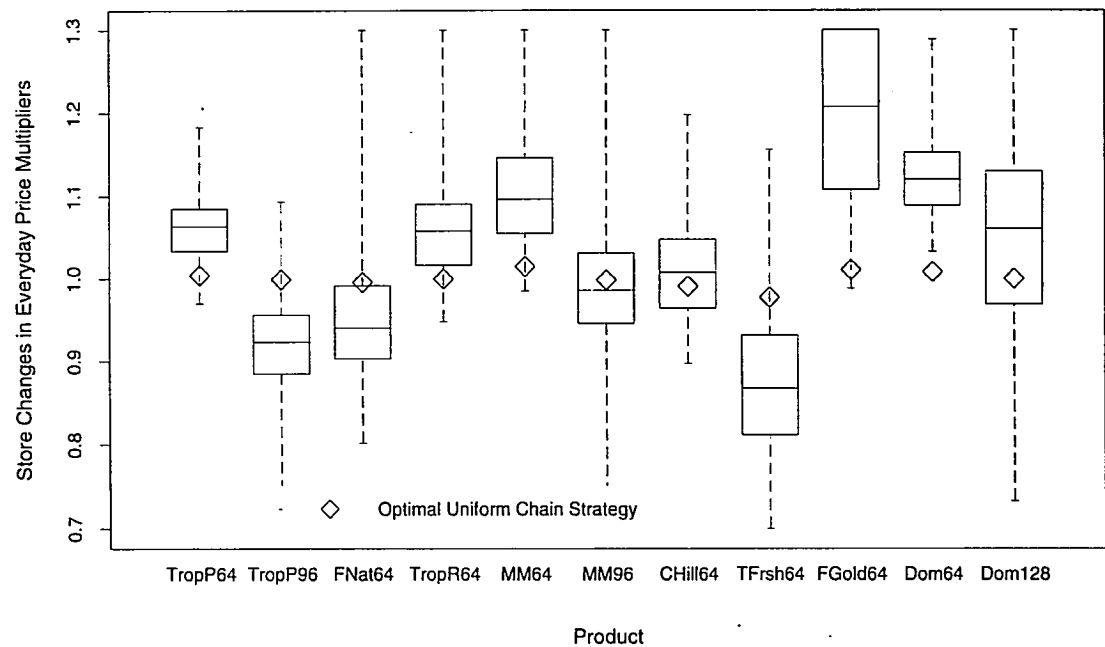
To represent the effects of these micro-marketing pricing changes, we overlay boxplots of the price changes across the stores for each brand in Figure 3. The variation of each boxplot around the diamond shows the differences that result from using store-level optimizations instead of an optimal chain-level price. As an example, consider the pricing effects on Tropicana Regular 64 oz. For this brand, the optimal chain price is virtually unchanged; however, in a micro-marketing strategy, it would increase in about 75 percent of the stores, and decrease in the remainder. Of course, the average price constraint means that any price increases in Tropicana Regular 64 oz. must be offset by price decreases of other products within the same store or shifting consumers to less expensive products. Certain products have very broad ranges in price movements, like Dominicks 128 oz., where some stores may have 30 percent price increases and others stores have 20 percent price decreases.

As a final step, we can further relax the constraints in (3.4) and (3.5) by holding average prices and revenues at the chain-level instead of the store-level:

$$\bar{p} = \sum_s w_s \bar{p}_{sr} \quad x = \sum_s x_{sr} \quad (3.8)$$

where $w_s = x_s/x$, x denotes the chain revenues, and \bar{p} denotes average chain price. This new micro-marketing strategy with the constraints enforced at a

Figure 3 Boxplots of Constrained Optimal Price Changes Across Stores.



chain-level (line 3 in Table 5) results in a 4.5 percent increase in expected profits, or an additional .6 percent increase over the solution with the constraints at a store-level. The choice between enforcing these constraints at a chain-level instead of a store-level reflects whether the retailer believes its image is formed at a store- or chain-level.

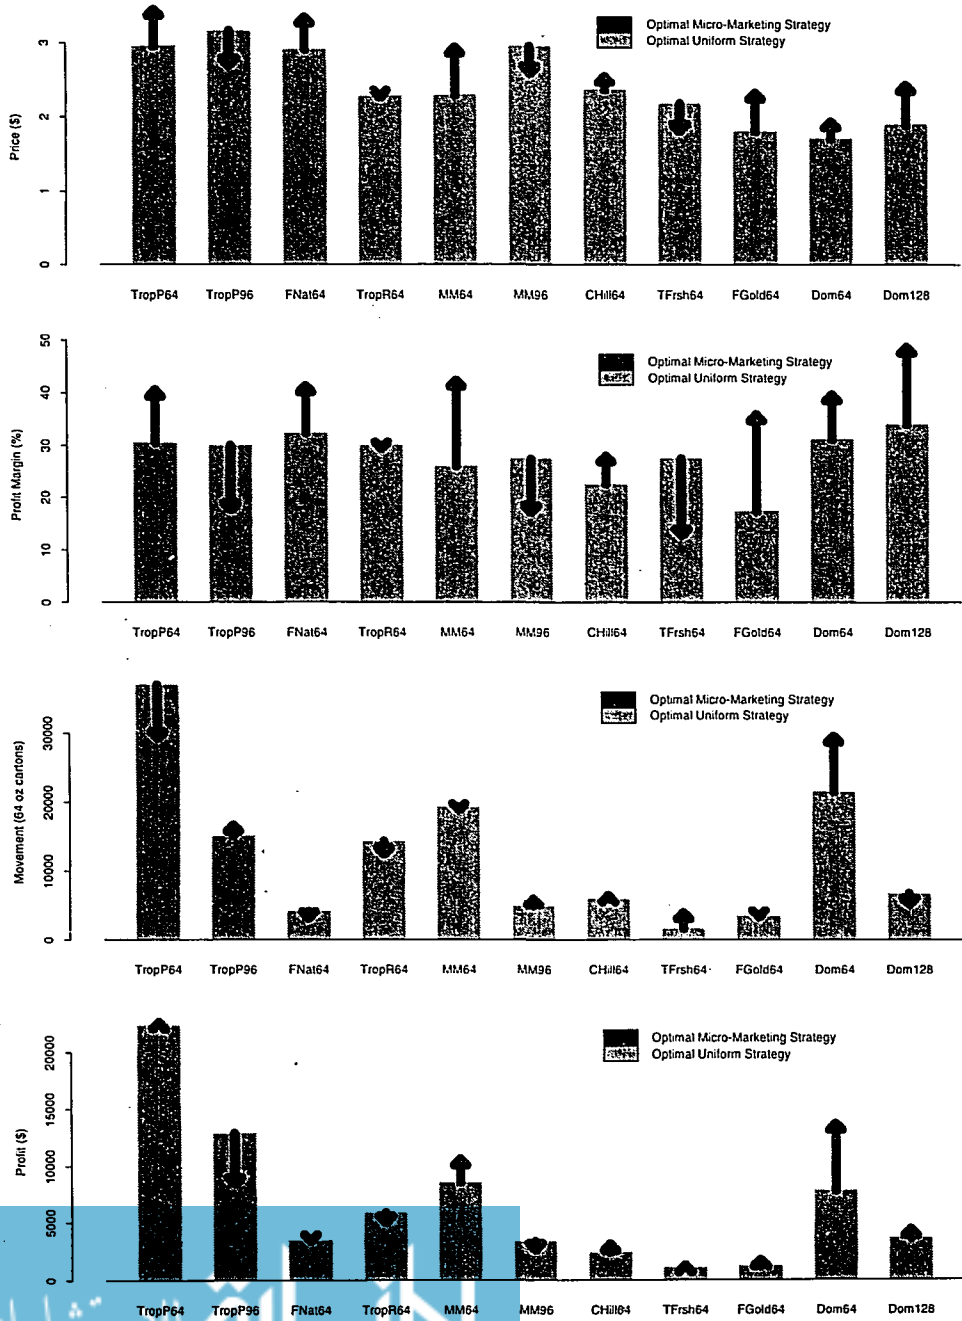
Given the small magnitude of some of these profit increases, it is worthwhile to ask whether these changes are subject to large estimation errors, i.e., statistically significant. To answer this question, we evaluate the probability that these new strategies increase expected profits over a uniform strategy, $\Pr[E[\pi_{new}] > E[\pi_{uniform}]]$. (See the appendix for a discussion of estimating these probabilities using the Gibbs sampler.) The probability that the expected profits from these new optimal strategies (lines 1, 2, and 3) exceed those of the current uniform strategy is greater than .99 for all these strategies (see the last column in Table 5). Alternately, we can compute the probability that the expected profits from these new micro-marketing

strategies (lines 2 and 3) exceed those of an optimal chain strategy (line 1), again the estimated probabilities exceed .99. The conclusion is that we can have a high degree of confidence that expected profits would increase using our new micro-marketing strategies.

3.4. Examining the Effects of Micro-Marketing for a Selected Store

To better illustrate how this proposed micro-marketing strategy works we select one store from an upscale northern suburb of Chicago and plot the price, movement, gross profit margins, and profits for each brand in Figure 4. The bars in this figure represent the values of each of these variables under the optimal uniform strategy as described above. The first panel of Figure 4 shows the median price levels for each brand. The arrows indicate the direction and magnitude of any changes as prescribed by the optimal micro-marketing strategy as previously discussed (line 2 of Table 5). First, notice that there are quite a few fairly large price increases for Tropicana Premium 64 oz.,

Figure 4 Effects of Micro-marketing for Price, Movement, and Profits for a Selected Store.



Florida's Natural 64 oz., Minute Maid 64 oz., Florida Gold 64 oz., and Dominicks 128 oz. In addition, there are several sizeable price decreases for Tropicana Premium 96 oz., Minute Maid 96 oz., and Tree Fresh 64 oz.

The purpose of these price changes is to encourage consumers to substitute to different products. This particular store is less price sensitive than the average store, and if possible the retailer would wish to increase all its prices. However, the average price and revenue constraints force the retailer to hold these variables at their current levels. Therefore, the pricing changes need to be engineered in such a way as to exploit store-level differences in substitution and encourage consumers to switch to more profitable bundles of products without large changes in existing demand. For this particular store, the asymmetry between the premium and store brands is more pronounced. Since the wholesale costs for these products are quite different the retailer will try to encourage substitution to more profitable bundles of products, as depicted by the changes in profit margins in the second panel of Figure 4.

The price increase of Tropicana Premium 64 oz. results in a primary decrease in demand, as illustrated by the lower quantity movement in the third panel of Figure 4. As a secondary effect, some consumers will switch to the lower quality national brands; this offsets the reduced movement that would have resulted from price increases for several national brands. In addition, some of the increase in the national brands will induce some consumers to switch to the lower quality store brands. It is this substitution that explains why the sales of Dominick's 64 oz. actually increases, even though there was a small increase in its price.⁶ The reason that these price changes cannot be implemented at a chain-level is that these asymmetric substitution effects vary across the stores. The average price constraint is satisfied since the share of the lower priced

⁶To verify that the large negative cross-price elasticities of Tropicana Premium 96 oz. and Dominicks 128 oz. are not driving these results, we repeat this analysis and hold their prices constant. Our projected profits increases will be approximately 26 percent less than the case where they are allowed to change (specifically, line 2 of Table 5 would be 2.8 percent). But this decrease is roughly in line with the fact that these brands have above average profit margins and account for 18 percent of category sales.

store brands has increased and the share of the higher priced national brands has decreased. Finally, profits have increased since the retailer is selling a great deal more of the store brands as shown in the fourth panel of Figure 4.

4. Improving the Retailer's Current Zone-Pricing Strategy

A common retailing pricing strategy is to cluster stores and then proportionately raise or lower prices in each cluster. These zone-pricing strategies represent a fairly limited form of micro-marketing, since they do not take advantage of differential brand effects. In addition, they may result in overall changes in average prices and revenues unlike our constrained pricing strategies presented in § 3.

Currently, DFF follows a zone strategy in which they segment their stores into high, medium, and low price zones. This segmentation is largely decided by competitive characteristics of the store's trading area. The prices of those stores close to warehouse competitors are lowered by 10 percent, and those stores in urban locations are increased by 10 percent. For example, a store assigned to the high price zone results in all prices being increased by 10 percent, i.e., $p_{1s} = p_{2s} = \dots = p_{11,s} = 1.1$. Table 6 lists the posterior mean and standard deviation of expected profits for DFF's current zone strategy. As before, we employ a uniform strategy as a basis for comparison (gross yearly profits for this category are around \$3.3mm). DFF's current zone strategy increases expected profits by .4 percent over a uniform chain strategy (line 1 of Table 6).

A natural question is to consider whether the retailer's current assignment of stores to each pricing zone can be improved. The best mechanism for assigning stores would be to evaluate their profit function and then assign stores with the maximal profit increases (or largest store profit gradient) to the high price zone and the stores with minimal profit increases to the low price zone. To keep the current zone strategy comparable with this new strategy, we allocate the same number of stores to each price zone. The expected profit increase using these optimal zones is 3.4 percent (line 2 of Table 6). In nominal terms, this represents an eight-fold increase in profits that are now

Table 6 Expected Gross Profit Changes Under Various Zone-Pricing Strategies

Description of Pricing Strategy	Relative to a Uniform Chain Pricing Strategy			
	Expected Profits	Expected Increase	% Change in Expected Profits	Prob[Expected Increase >0]
Uniform Prices across all Stores	\$3,330,900 (11,900)			
1 Current Zone Strategy	\$3,344,100 (11,400)	+\$13,200 (7,300)	+ .4% (.2)	.97
2 Zones using Profit Gradients	\$3,442,900 (17,900)	+\$112,000 (9,000)	+3.4% (.3)	1.00
3 Optimal Micro-Marketing Strategy with Zones	\$3,662,500 (22,700)	+\$331,600 (20,600)	+10.0% (.6)	1.00

Note. The standard deviations of the posterior are given in parentheses below the posterior means.

attributable to micro-marketing policies using DFF's current zone strategy. The reason for the superior performance over the retailer's current zone assignment is that the retailer relies almost solely upon competitive characteristics without considering the store's history of price response and the demographic profile of the store's trading area.

One final suggestion would be to combine the optimal pricing strategies at a store-level from § 3 with these optimal price zones. That is, we assign stores to high, medium, or low price zones based upon their profit gradients. Then we optimize the prices within each store subject to the average price constraint as before. For example, suppose a store is assigned to a high price zone with 10 percent higher price levels, then prices are optimized subject to (3.4) and (3.5) with a 10 percent higher average price. This new strategy yields almost a 25-fold increase over the gains of the current uniform strategy or a 10 percent increase in expected profits (see line 3 of Table 6). The basic difference between the line 3 of Table 5 and line 3 of Table 6 is whether the average price and revenue conditions must strictly hold at the chain-level. Line 3 of Table 6 allows the chain average price and revenues to increase slightly over current levels, while line 3 of Table 5 holds them strictly at their current level. These differences can be explained by the inelasticity of this category's aggregate demand, which seems to indicate overall price increases are profitable.

5. The Relationship Between Modeling and Pricing Decisions

Typically, modeling and pricing decisions are treated independently of one another. However, the choices a modeler makes have important consequences on the pricing strategies derived. A key decision in our hierarchical Bayesian model is the prior that assesses the commonalities across stores, i.e., how much shrinkage should occur. In our problem, the amount of shrinkage is largely influenced by the prior on V_p . We have parameterized this prior such that the mean of its inverse is scaled by a single parameter, k (see the appendix for further discussion). If $k = 0$ and there are no demographic effects ($\gamma = 0$), then we have a pooled model, if $k = 1$ we have an empirical Bayes prior, and if k is large (which for our case is around 10) our estimates are close to the individual LS store models.

We choose a value of $k = 1$ to reveal our prior beliefs that shrinkage is helpful and to reflect the fact that pooled models tend to perform well compared with individual store models. To validate the choice of our prior and better understand its influence on the posterior, we compute the posterior odds for several values of k . Based upon the full sample, we find $k = 1$ to be 22 times more likely than the model with $k = .1$ in terms of the log of the posterior odds ratio, and 52 times more likely than one with $k = 5$. Clearly in natural units the posterior odds strongly favor $k = 1$.

Essentially the data is indicating that individual LS models over-fit the data. However, a pooled model or even those with strong priors that induce strong shrinkage are too severe. This result can also be seen from the likelihood ratio test reported in § 2.1. The reason that the hierarchical model with a moderate amount of shrinkage is preferred is that it allows for estimates that fall between these two extremes.

An alternative validation exercise is to compare the MSE and MAD (mean absolute deviation) of out-of-sample predictions for these various models. The out-of-sample predictions are computed by dividing the sample into two halves, each with roughly 60 weeks, with the first half used solely for estimation. The MSE and MAD statistics are provided for various models in Table 7 and are listed relative to the absolute value of the individual LS models given in the last row. The LS estimates by definition minimize the in-sample MSE, so it is not surprising that the in-sample MSE of the Bayes model exceeds those of the individual LS models. However, the Bayes models show improved out-of-sample predictions. The Bayes model with a weak prior shows a 9 percent reduction in the MSE compared with the individual LS models. As this prior is strengthened, we can observe a decrease in the out-of-sample MSE by about 20 percent. We can note a similar pattern in the out-of-sample MAD, although the maximum gain indicated by this statistic is 10 percent.

The improvement in predictive ability for the Bayes models comes from the tradeoff in bias and variance of the parameter estimates. These results show that the predictions can be improved by borrowing information across stores in a hierarchical Bayesian model.

While predictive validation cannot prove that a particular prior is correct, it shows that strong or moderate priors in commonalities of cross-store parameter variation is supported by the data. In this paper we will not try to resolve the question about which prior is correct, although clearly the data favor moderate or strong amounts of shrinkage. Instead, we suggest the analyst follow a Bayesian approach to reflect his prior beliefs about heterogeneity in price response across stores. For further discussion of this problem, we refer the interested reader to Montgomery and Rossi (1997), who consider more flexible prior specifications that yield greater insights into tradeoffs between heterogeneity and model specification.

Better out-of-sample predictions are desirable, but our real objective is to optimize profits using a model that reflects the current environment. To understand the impact of these prior assessments about store commonalities, we list the percentage change in expected profits for the various pricing strategies that were discussed in Table 5 from § 3.3 and recompute them for strong, moderate, and weak priors in Table 8. Notice the second column, labeled "Moderate Prior on Store Commonalities," is the same as in Table 5. For the optimal micro-marketing strategy (line 2) that we presented earlier we expect profits to increase from 3.9 percent to 4.9 percent, a 25 percent increase, using a weak prior instead of a moderate prior. When using an optimal micro-marketing strategy that keeps revenue and average price at the chain-level constant (line 3) we observe a increase of similar magnitude, from a 4.5 percent increase in expected profits to a 5.9 percent increase.

Table 7 Sensitivity of Out-of-Sample Predictions to Prior on Store Commonalities

Model	MSE Relative to LS Models		MAD Relative to LS Models	
	In-sample	Out-of-sample	In-sample	Out-of-sample
Individual LS Models	1.00	1.00	1.00	1.00
Bayes, Weak Prior ($k = 5$)	1.02	0.91	1.01	0.96
Bayes, Moderate Prior ($k = 1$)	1.07	0.82	1.02	0.91
Bayes, Strong Prior ($k = .1$)	1.18	0.80	1.07	0.90
Pooled Chain Model	2.13	1.14	1.58	1.12
Base MSE of LS Models	0.17	0.41	0.29	0.46

Table 8 Sensitivity of % Change in Expected Profits to the Prior on Store Commonalities

Description of Pricing Strategy	Prior on Store Commonalities		
	Strong	Moderate	Weak
1 Optimal Uniform Strategy	+1.8% (.3)	+ .4% (.1)	+ .5% (.1)
2 Optimal Micro-Marketing Strategy	+3.2% (.5)	+3.9% (.6)	+4.9% (.9)
3 Optimal Micro-Marketing Strategy with Constraints at the Chain-Level	+3.6% (.5)	+4.5% (.6)	+5.9% (1.0)

Note. The standard deviations of the posteriors are given in parentheses below the posterior means.

This means that micro-marketing becomes more important if we assume less a priori about store commonalities. If a researcher is only willing to make weaker assumptions about store commonalities, as with individual store models, the expected profit gain from micro-marketing is larger. However, the concern of individual store models is that they over-fit the data. What is driving the increased profits attributable to micro-marketing in these weaker priors is greater variation in store-level demand. As the prior is weakened, larger store-level demand differences occur, which translate into larger gains for micro-marketing. It should be noted that the variance of expected profits also increases since individual store demand models become more difficult to estimate. Regardless of the strength of our prior beliefs, we still reach the same substantive conclusions that store-level differences in price sensitivity can be measured and translated into profitable micro-marketing strategies.

6. Conclusions

This paper has shown that store differences in demand at the product-level can be measured and translated into micro-marketing pricing strategies that result in significant expected profit gains for the retailer. The difficult implementation task of micro-marketing (Blattberg 1988) has been addressed through the development of constrained pricing strategies that allow the retailer to retain their current store image while still

making store-level changes in everyday prices. Essentially we have separated the micro-marketing component of the problem from creating a base pricing and promotional strategy for the chain. This allows the retailer to retain its current managerial practices and positioning strategies, while at the same time exploiting store-level differences in demand without effecting the store's image. A desirable aspect of our approach is that it leverages existing datasets to measure store-level demand, so that additional and costly new datasets are not needed to implement micro-marketing strategies.

We have demonstrated the statistical significance of these expected profit gains, but there is also the question of the managerial significance of these micro-marketing profit gains. Our results predict that micro-marketing strategies increase gross profits anywhere from 3.9 percent to 10 percent over a uniform chain pricing strategy, depending upon the form of the micro-marketing strategy. Given DFF's willingness to engage in a zone-pricing strategy that yields a .4 percent increase in expected gross profits, a 10- to 25-fold increase in profits attributable to micro-marketing represents a sizeable gain.

Gross profits only consider the difference between the retailer's price and cost, and they do not reflect administrative and other selling costs. For refrigerated juices, gross profit margins are around 25 percent, which is the average for supermarket retailers (Supermarket Business 1992). On the other hand, operating profit margins were around 3 percent in 1993 (see publicly disclosed financial statements and financial summary in Forbes 1994). If we assume that these micro-marketing increases are consistent across categories,⁷ then a 4 percent increase in gross profits translates into

⁷We are optimistic that these results generalize to other categories and retailers, since previous research (Hoch et al. 1995) suggests similar variation in demand across stores. In addition, the supermarket retailer considered in this paper follows a Hi-Lo pricing strategy, which is the most commonly employed pricing strategy in this industry, and operates in a large metropolitan area with a diverse population. However, wide-spread application of these micro-marketing strategies may result in unforeseen difficulties, perhaps due to added costs or dramatic consumer reactions. See the discussion later in the conclusions where we consider the limitations of this study.

a much larger increase in operating profits of 33 percent. A 10 percent increase in gross profits would almost double operating profits with an increase of 83 percent. In nominal terms these increases would represent gains in the tens of millions of dollars in operating profits to a regional chain. Certainly these are sizeable profit gains for any retailer.

These predicted profit increases would not be possible without the flexibility of product-level price changes in each store. The retailer's current zone pricing strategies are not able to gain such a fine pricing result, since zone pricing uses overall price increases or decreases at a store-level. Instead, our micro-marketing approach caters to local purchase patterns by measuring unique price response profiles for each store. By considering each store's profile individually, we have been able to look for price changes that increase profits by encouraging substitution between similar types of bundles of products with the same sales dollar value but higher profits. This substitution to alternate bundles is achieved without changing average price or losing market share to competing retailers, but relies exclusively on individual store differences in price response. This overcomes a criticism of current pricing research (Hoch, Dreze, and Purk 1994) that requires chains to reposition themselves with higher overall prices, since under our micro-marketing strategy image is unchanged.

There are several limitations to our study that deserve further attention. First, the profit increases given in this paper are all expected gains. In practice, retailers may experience higher or lower returns due to random effects from uncontrolled factors. In addition, these profit predictions are made conditional upon our model specification. Any misspecifications in the model or additional uncertainty about the functional form of demand (Kalyanam 1996, Montgomery and Bradlow 1997) may alter these predictions.

Another limitation is that we have only considered a partial solution of the pricing problem that has limited its attention to everyday pricing. There are many marketing mix elements, like promotional prices and assortments, that can be manipulated at a store-level. Previous research into promotions by Jeuland and Narasimhan (1985) advanced a price discrimination mechanism for promotions that would suggest that

micro-marketing feature policies could be successful. If we allow feature prices to be store specific, we would expect to double our expected profit gain (Montgomery 1994). Since these results do not fully exploit the retailer's dynamic cost structures (i.e., forward buying and promotional offers), we would expect these gains might even be greater due to more efficient allocation of inventory and promotional budgets. However, our information set does not allow us to properly gauge consumer and competitive reaction to promotional changes, so we leave the joint specification of micro-marketing pricing and promotion strategies to future research.

This leads us to reconsider whether our constraints to maintain price image are adequate, or essentially whether micro-marketing can be addressed with this limited information set. Even though our dataset is huge, it does not include price and promotional activity of competing retailers. We have controlled for this by attempting to preserve the current image. However, any misspecification of the store's image could result in poor pricing predictions. Under certain forms of consumer heterogeneity these constraints may be inadequate. Specifically, they may need to be extended to reflect other bounds (e.g., no price changes larger than 10 percent) or other constraints may need to be added (e.g., the average price of typical product bundles purchased by various consumer segments within the store should not change).

In addition, our reliance upon constrained solutions to approximate current competitive conditions may not be valid in all operating contexts. Specifically, they may "awaken" competition between retailers by bringing pricing into center-focus. This may lead competitors to abrupt marketing changes that are not foreseen by our model. Alternately, managers may wish to use micro-marketing to reposition a store and deliberately provoke competitive responses. Again, a solution to both these problems would be to directly include competitive information and reactions into the model, although this would make micro-marketing more difficult and require a more extensive model and dataset. Clearly, further research is needed on how a store's image is determined. As the components of store image are better identified, our hope is that the proposed

constraints can be augmented to ensure store image is properly measured.

Most of our price changes are relatively small in magnitude (less than 25 ¢), so changes in utility at the individual consumer level should also be small, although even these small changes may become problematic if micro-marketing strategies are implemented across multiple categories and the changes in individual utility levels are highly correlated across categories. If this is true, some consumers may begin to realize sizeable gains/losses and may have enough of an incentive to switch stores. Current research has found only low correlations in the price sensitivity of consumers across categories (Kim and Srinivasan 1995). However, this research question has not yet been resolved.

Finally, there may be circumstances under which consumers will react against differential store pricing strategies, which could result in structural changes of demand. If a consumer lives on a store's border region and experiences different prices for the same product within a chain, that consumer may become angry and switch retailers. However, the best argument against any adverse consumer reaction is that retailers currently have different prices for the same product in neighboring stores using zone pricing and adverse consumer reactions have not been observed. The broader ethical and public policy considerations of micro-marketing are beyond the scope of this paper, but are interesting areas for future studies.⁹

Appendix A: Prior Specifications and Estimation of Our Hierarchical Model

Besides the data the analyst must also supply the following parameters:

$$v_{\Sigma}, \bar{V}_{\Sigma}^{-1}, v_{\beta}, \bar{V}_{\beta}^{-1}, \bar{\theta}, V_{\theta} \quad (\text{A.1})$$

⁹The author would like to thank Peter Rossi and Steve Hoch for their valuable input. He also thanks Mark Bergen, Pete Fader, Kris Helsen, Rob McCulloch, Jagmohan Raju, and George Tiao for their comments, Dominick's Finer Foods, Information Resources Inc., and Market Metrics for their assistance and provision of data, and Xavier Dreze and Mary Purk for their indispensable help throughout. Finally, the author wishes to thank the editor, area editor, and two anonymous referees for their valuable comments. Financial support for this work was provided by the Micro-Marketing Project at the Graduate School of Business, University of Chicago.

Our priors on the error covariance matrices, Σ_s , and mean of the hyper-distribution, θ , are chosen to be diffuse relative to the sample:

$$\bar{V}_{\Sigma} = .1 I, v_{\Sigma} = 1, \bar{\theta} = 0, V_{\theta} = 10^6 I. \quad (\text{A.2})$$

The most crucial prior will be on V_{β} , which reflects the strength of the commonalities across the stores. Notice that in our case the number of stores is less than the dimension of V_{β} , therefore, to form a proper posterior distribution, we need to have an informative prior. The motivation of the parameterization of our prior on V_{β} is to shrink our parameter estimates somewhere between the pooled and individual LS estimates. We set \bar{V}_{β} to a diagonal matrix. To allow for proper scaling of the different coefficients, we set the diagonal elements equal to the product of the variance of the least squares estimates from the individual store models, s_i^2 , and a scaling parameter, k_i :

$$\bar{V}_{\beta} = v_{\beta} \bar{V}_{\beta}^* = v_{\beta} \begin{bmatrix} k_1^2 s_1^2 & & & \\ & k_2^2 s_2^2 & & \\ & & \ddots & \\ & & & k_p^2 s_p^2 \end{bmatrix}, \quad (\text{A.3})$$

where p is the dimension of β_i (or 154, i.e., 11 equations with 14 parameters each). We also set $v_{\beta} = 160$.

The relationship between V_{β}^{-1} and \bar{V}_{β}^{-1} can be seen by examining the mean and covariance matrix of this prior distribution: $E[V_{\beta}^{-1} | v_{\beta}, \bar{V}_{\beta}^{-1}] = v_{\beta} \bar{V}_{\beta}^{-1} = \bar{V}_{\beta}^{*-1}$ and $\text{Var}[V_{\beta}^{-1}] = 2 / v_{\beta} \bar{V}_{\beta}^{*-1} \otimes \bar{V}_{\beta}^{-1}$. For example, if $k_i = .1$, then our prior states that the expected standard deviation of store-specific random variation of our estimates will be 10 percent of those of the LS estimates. (Note that parameter variation is also be induced by the demographics.)

We are primarily concerned with parameter variation across stores and do not try to shrink the parameters within a store (i.e., across products) closer to each other. Therefore, we let each store and brand have their own intercepts. Also, to avoid a great deal of shrinkage in the constants, we set the scaling parameters of these parameters to one or k_i , whichever is greater. For simplicity, we set the scaling parameters for the other parameters equal, $k_i = \max(1, k)$ for constants and $k_i = k$ for all other parameters. We choose a value of $k = 1$ to reflect moderate commonalities a priori. For a further discussion of the effects of k , see § 5.

The nested structure of the model simplifies the specification of the hierarchical model, although estimation presents another problem. Even with natural conjugate priors an exact solution of the posterior distribution is not known.¹⁰ Therefore, we make use of a new technique in computational statistics, known as the Gibbs sampler, to estimate the marginal posterior distributions. For a good introduction to the Gibbs sampler, see Casella and George (1992). The

¹⁰The analytical solution that integrates Σ_s out of the joint distribution of β to derive the posterior distribution is not known. The difficulty arises as a result of the Wishart priors on Σ_s and V_{β} . To understand this problem, we refer the reader to the simpler case of trying to solve a single stage SUR model (Zellner 1971, pp. 240-246) for which the analytical solution is not known either.

Gibbs sampler requires randomly sampling from each of the conditional distributions sequentially.

An added benefit of the Gibbs draws is that we may compute an estimate of the marginal posterior distribution of the expected profit function, which incorporates the uncertainty of the parameter estimates. Traditionally, the posterior means of the parameter estimates are substituted into the profit function. Blattberg and George (1992) show that this method does not lead to an optimal pricing solution due to the nonlinearity of the profit function. A detailed technical report is available upon request from the author with a complete description of the algorithm.

References

Allenby, Greg M. (1989), "A Unified Approach to Identifying, Estimating and Testing Demand Structures with Aggregate Scanner Data," *Marketing Science*, 8, 265-280.

Armstrong, Marcia K. (1991), *Retailer Response to Trade Promotions: An Incremental Analysis of Forward Buying and Retail Promotion*, Unpublished Ph.D. Dissertation, University of Texas at Dallas.

Bell, David R. (1995), *Consumer Response to Retailer Pricing Strategies*, Unpublished Ph.D. Dissertation, Graduate School of Business, Stanford University, Stanford, CA.

Blattberg, Robert C. (1988), "Assessing and Capturing the Soft Benefits of Scanning," *Coca-Cola Retailing Research Council Study*.

— and Edward I. George (1991), "Shrinkage Estimation of Price and Promotional Elasticities: Seemingly Unrelated Equations," *JASA*, 86, 304-315.

— and — (1992), "Estimation Under Profit-Driven Loss Functions," *Journal of Business and Economic Statistics*, 10, 437-444.

— and Scott A. Neslin (1990), *Sales Promotion: Concepts, Methods, and Strategies*, NJ: Prentice-Hall: Englewood Cliffs.

— and Kenneth J. Wisniewski (1989), "Price-Induced Patterns of Competition," *Marketing Science*, 8, 291-309.

Bodapati, Anand V. (1996), *The Impact of Out-of-Store Advertising on Store Sales*, Unpublished Ph.D. Dissertation, Graduate School of Business, Stanford University, Stanford, CA.

Casella, George and Edward I. George (1992), "Explaining the Gibbs Sampler," *American Statistician* 46, 167-174.

Cutler, Kirstie and Christopher Rowe (1990), "Scanning in the Supermarket: For Better or Worse? A Case Study in Introducing Electronic Point of Sale," *Behavior and Information Technology*, 9, 157-169.

Deaton, Angus and John Muellbauer (1980), *Economics and Consumer Behavior*, Cambridge UK: Cambridge University Press.

Dickson, Peter R. and Alan G. Sawyer (1990), "The Price Knowledge and Search of Supermarket Shoppers," *Journal of Marketing*, 54, 42-53.

Dreze, Xavier (1995), *Loss Leaders: Store Traffic and Cherry Picking*, Unpublished Ph.D. Dissertation, Graduate School of Business, University of Chicago, Chicago, IL.

Forbes, "Annual Report on American Industry," January 3, 1994.

Garland, Ron (1992), "Pricing Errors in the Supermarket: Who Pays?" *International Journal of Retail and Distribution Management*, 20, 25-30.

Gelfand, Alan E., Susan E. Hills, Amy Racine-Poon, and Adrian F. M. Smith (1990), "Illustration of Bayesian Inference in Normal Data Models Using Gibbs Sampling," *JASA* 85, 972-985.

— and Adrian F. M. Smith (1990), "Sampling-Based Approaches to Calculating Marginal Densities," *JASA* 85, 398-409.

Goodstein, Ronald C. (1994), "UPC Scanner Pricing Systems: Are They Accurate?" *Journal of Marketing*, 58, 20-30.

Gupta, Porter, and Wittink (1993), "Dangers in Using Market Level Data in Determining Promotion Effects," *Marketing Science Institute*, No. 93-115 Working Paper.

Hoch, Stephen J., Xavier Dreze, and Mary E. Purk (1994), "EDLP, Hi-Lo, and Margin Arithmetic," *Journal of Marketing*, 58, 16-27.

—, Byung-Do Kim, Alan L. Montgomery, and Peter E. Rossi (1995), "Determinants of Store-Level Price Elasticity," *Journal of Marketing Research*, 32, 17-29.

— and Mary E. Purk (1993), "Pricing Results and Recommendations from the Refrigerated Juice Category, Report from Phase II of the Micro-Marketing Project (Dated: April 6, 1993)," Graduate School of Business, University of Chicago, Chicago, IL.

Jeuland, Abel and Chakravarthi Narasimhan (1985), "Dealing-Temporary Price Cuts-by Seller as a Buyer Discrimination Mechanism," *Journal of Business*, 58, 295-308.

Kalyanam, Kirthi (1996), "Pricing Decisions Under Demand Uncertainty: A Bayesian Mixture Model Approach," *Marketing Science*, 15, 3, 207-221.

Kim, Byung-Do and Kannan Srinivasan (1995), "Investigation of Consumers' Price Sensitivity Across Multiple Product Categories: An Empirical Generalization," Working Paper, Carnegie Mellon University, Pittsburgh, PA.

Kumar, V. and Robert P. Leone (1988), "Measuring the Effect of Retail Store Promotions on Brand and Store Substitution," *Journal of Marketing Research*, 25, 178-85.

Lindley, David V. and Adrian F. M. Smith (1972), "Bayes Estimates for the Linear Model," *Journal of the Royal Statistical Society, Series B*, 34, 1-41.

Montgomery, Alan L. (1994), *The Impact of Micro-Marketing on Pricing Strategies*, Unpublished Ph.D. Dissertation, Graduate School of Business, University of Chicago, Chicago, IL.

— and — (1997), "Why Overconfidence About the Functional Form of Demand Models Can Lead to Overpricing," Working Paper, The Wharton School, University of Pennsylvania, Philadelphia, PA.

— and Peter E. Rossi (1997), "A Hierarchical Bayesian Approach to Estimating Price Elasticities," Working Paper, The Wharton School, University of Pennsylvania, Philadelphia, PA.

NAG (1991), *Mark 15 Fortran Libraries*, Downers Grove, IL: The Numerical Algorithms Groups, Inc.

Rossi, Peter E., Robert E. McCulloch, and Greg M. Allenby (1996), "The Value of Purchase History Data in Target Marketing," *Marketing Science*, 15, 4, 321-340.

Smith, Adrian F. M. (1973), "A General Bayesian Linear Model," *Journal of the Royal Statistical Society, Series B*, 35, 67-75.

Supermarket Business, September 1992.


Walters, Rockney G. (1991), "Assessing the Impact of Retailer Price Promotions on Product Substitution, Complementary Purchase, and Interstore Sales Displacement," *Journal of Marketing* 55, 17-28.

— and Scott B. MacKenzie (1988), "A Structural Equations Analysis of the Impact of Price Promotions on Store Performance," *Journal of Marketing Research*, 25, 51-63.

This paper was received August 23, 1995, and has been with the author 14 months for 3 revisions; processed by Greg M. Allenby

Wittink, Dick R., Michael J. Addona, William J. Hawkes, and John C. Porter (1988), "SCAN*PRO: The Estimation, Validation, and Use of Promotional Effects Based on Scanner Data," Working Paper, Johnson Graduate School of Management, Cornell University, Ithaca, NY.

Zellner, Arnold (1971), *An Introduction to Bayesian Inference in Econometrics*. New York: Wiley.



Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.